2022 AI HW Summit: The Annual Top 10 List

Karl Freund

Founder and Principal Analyst

Cambrian-AI Research LLC

Karl.Freund@Cambrian-ai.com

@karlfreund

The 2021 Top 10 List

➔ All are still Evolving

1. MLPerf Benchmarks
2. Synopsys DSO.ai
3. Sambanova, Groq, & Tenstorrent
4. AWS: Inferentia
5. Google: TPU-V4
6. Intel Habana Labs Gaudi
7. Graphcore 2nd Generation and the IPU-Machine
8. NVIDIA Grace
9. Cerebras WSE-2 and Brain-Scale AI
10. …. (The founding or Cambrian-AI, of course!)

Why rob Banks? That's where the money is.

300+ Trillion
Inferences per day

5B+
Machine translations per day

Source: 2020 Omdia

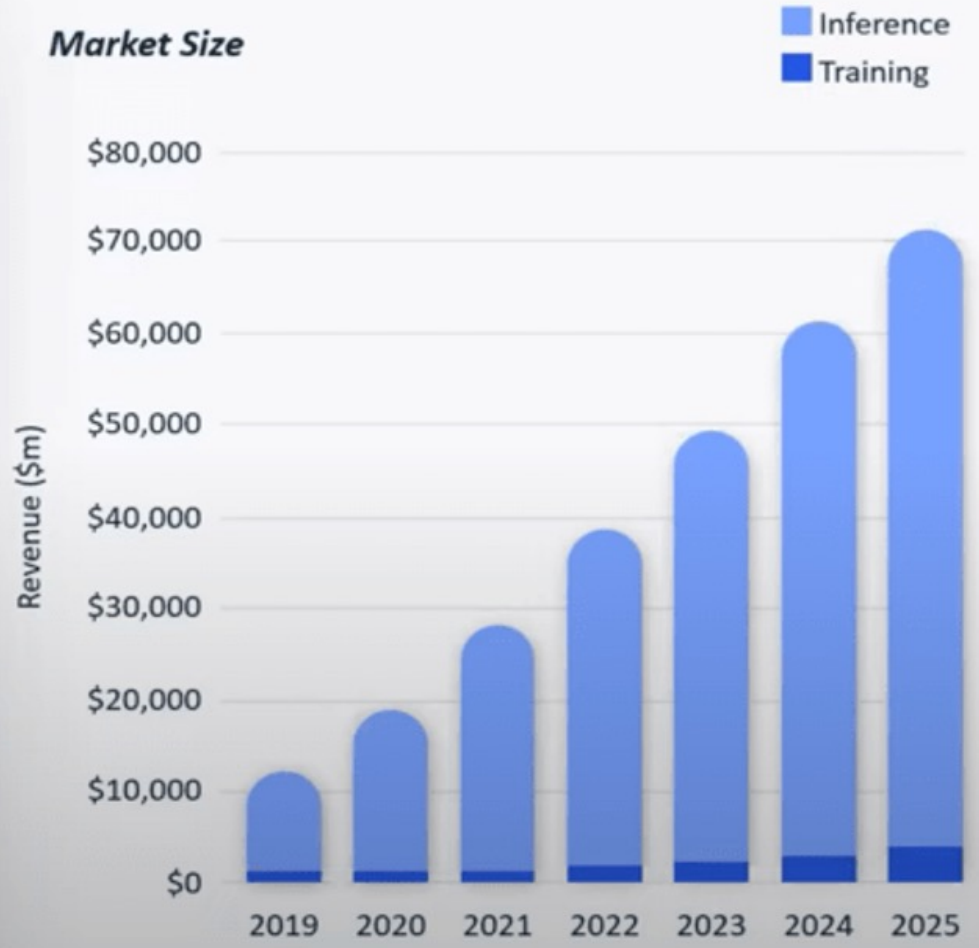

MLPerf Inference 2.1

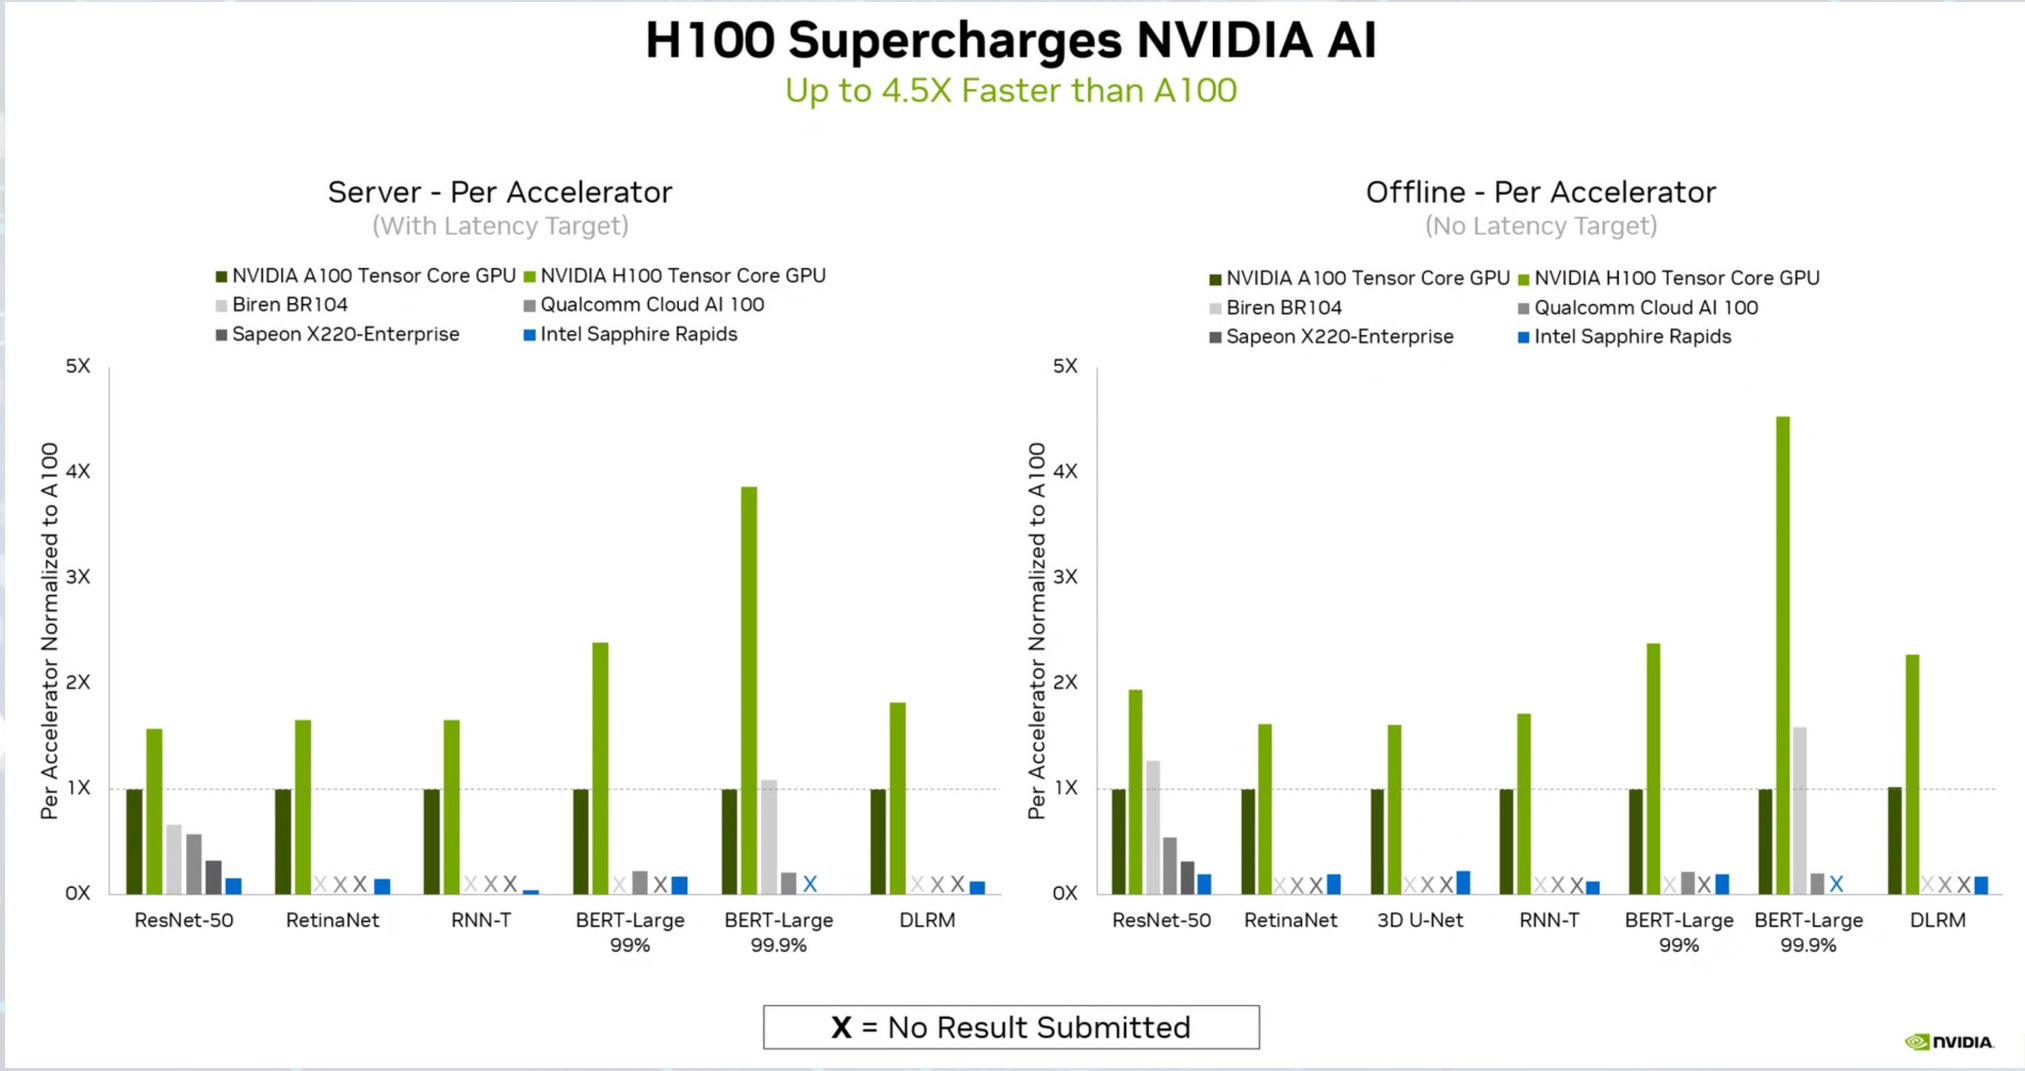

Qualcomm Leads MLPerf Efficiency

Over 200 submissions with Server Ecosystem

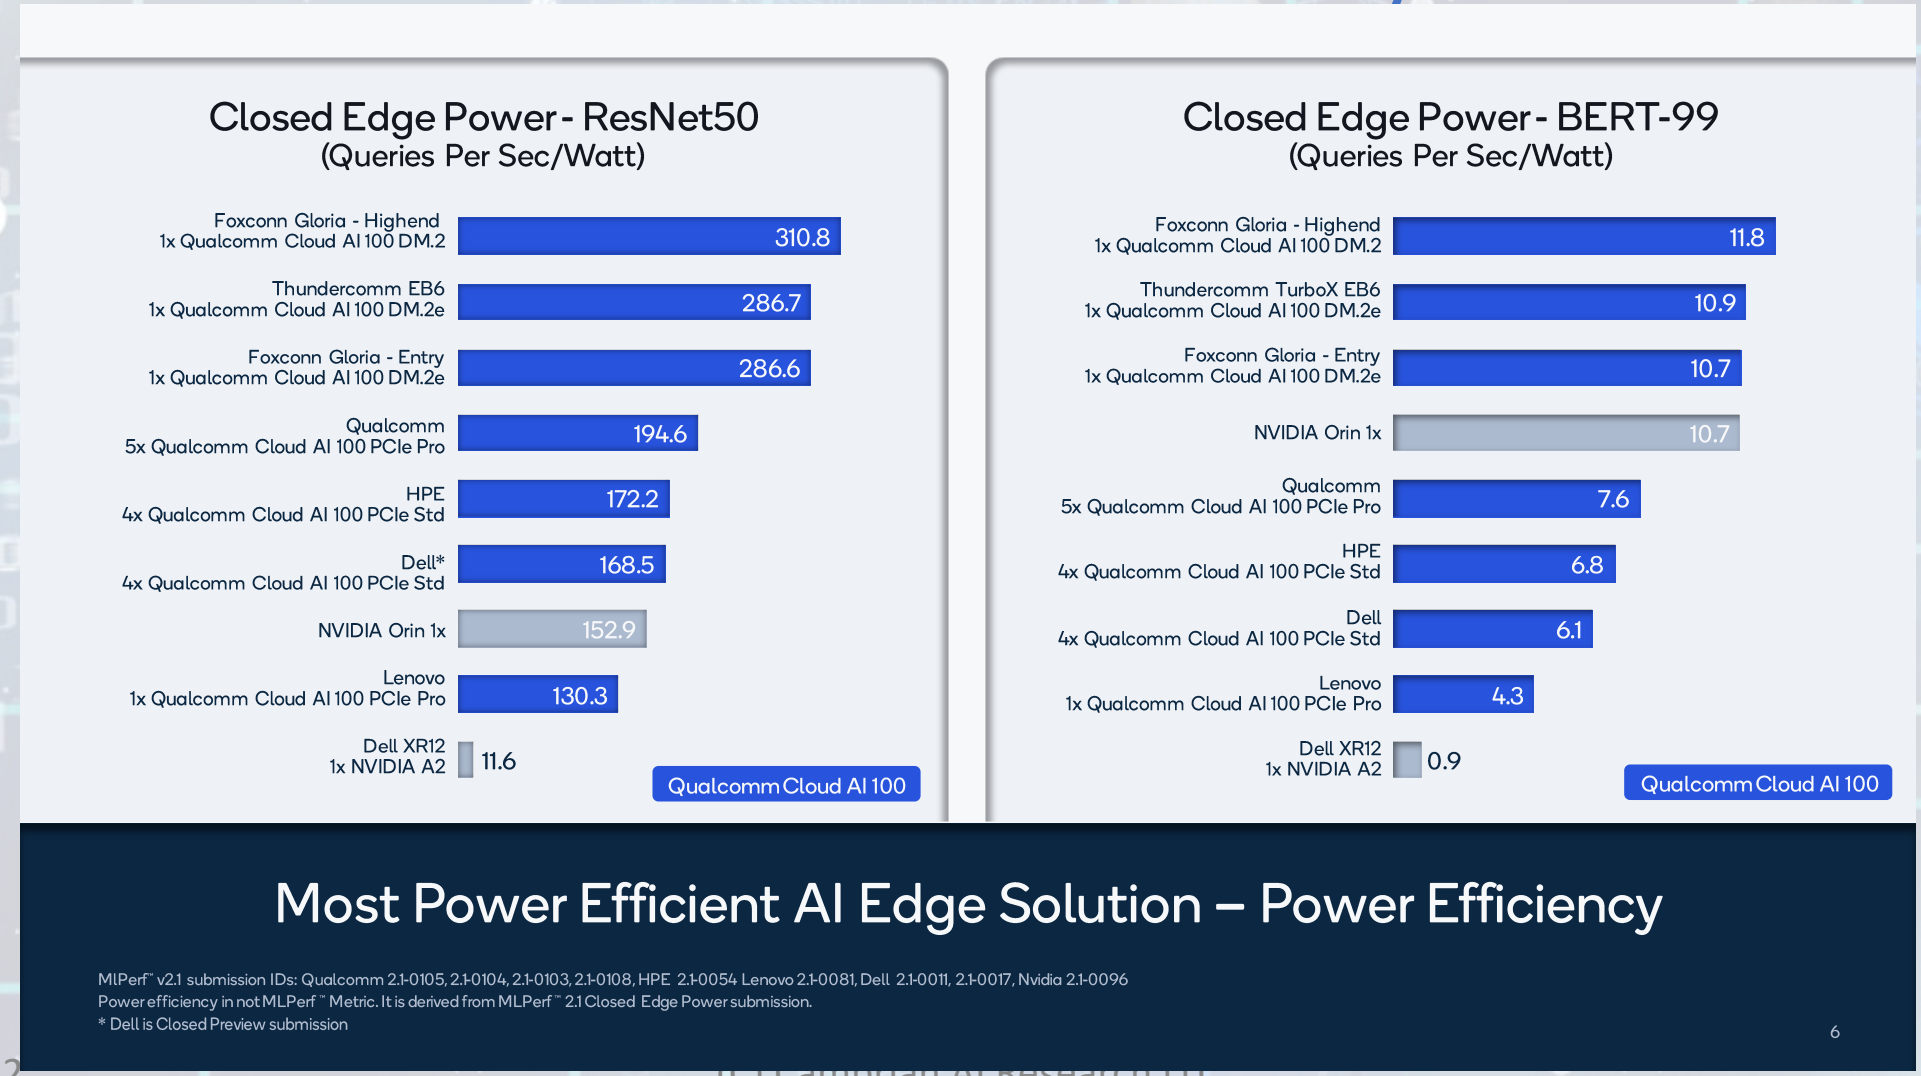

Most Power Efficient AI Edge Solution – Power Efficiency

MLPerf™ v2.1 submission IDs: Qualcomm 2.1-0105, 2.1-0104, 2.1-0103, 2.1-0108, HPE 2.1-0054 Lenovo 2.1-0081, Dell 2.1-0011, 2.1-0017, Nvidia 2.1-0096
Power efficiency in not MLPerf ™ Metric. It is derived from MLPerf ™ 2.1 Closed Edge Power submission.
* Dell is Closed Preview submission

The Top 10 AI HW Innovations of 2022

1. **GrAI Matter Labs** – High Fidelity Edge processing
2. **D-Matrix:** In-Memory Computation
3. **Untether.AI:** At-Memory Computation
4. **Mythic:** Here comes Analog compute
5. **Esperanto** : 1000 RISC-V cores
6. **SimaAI:** MLSoC provides SW-Centric ML
7. **AMD**: Massive FLOPS w/ A100-class TOPS
8. **Graphcore**: Wafer on Wafer, and Good Computer
9. **Intel Habana Labs**: Gaudi 2 doubles NVIDIA A100 throughput
10. **NVIDIA** Hopper and Grace **Superchips**: the Next Generation of Compute

GrAI Matter Labs:
High Fidelity real-time inference

D-Matrix: In-Memory Compute

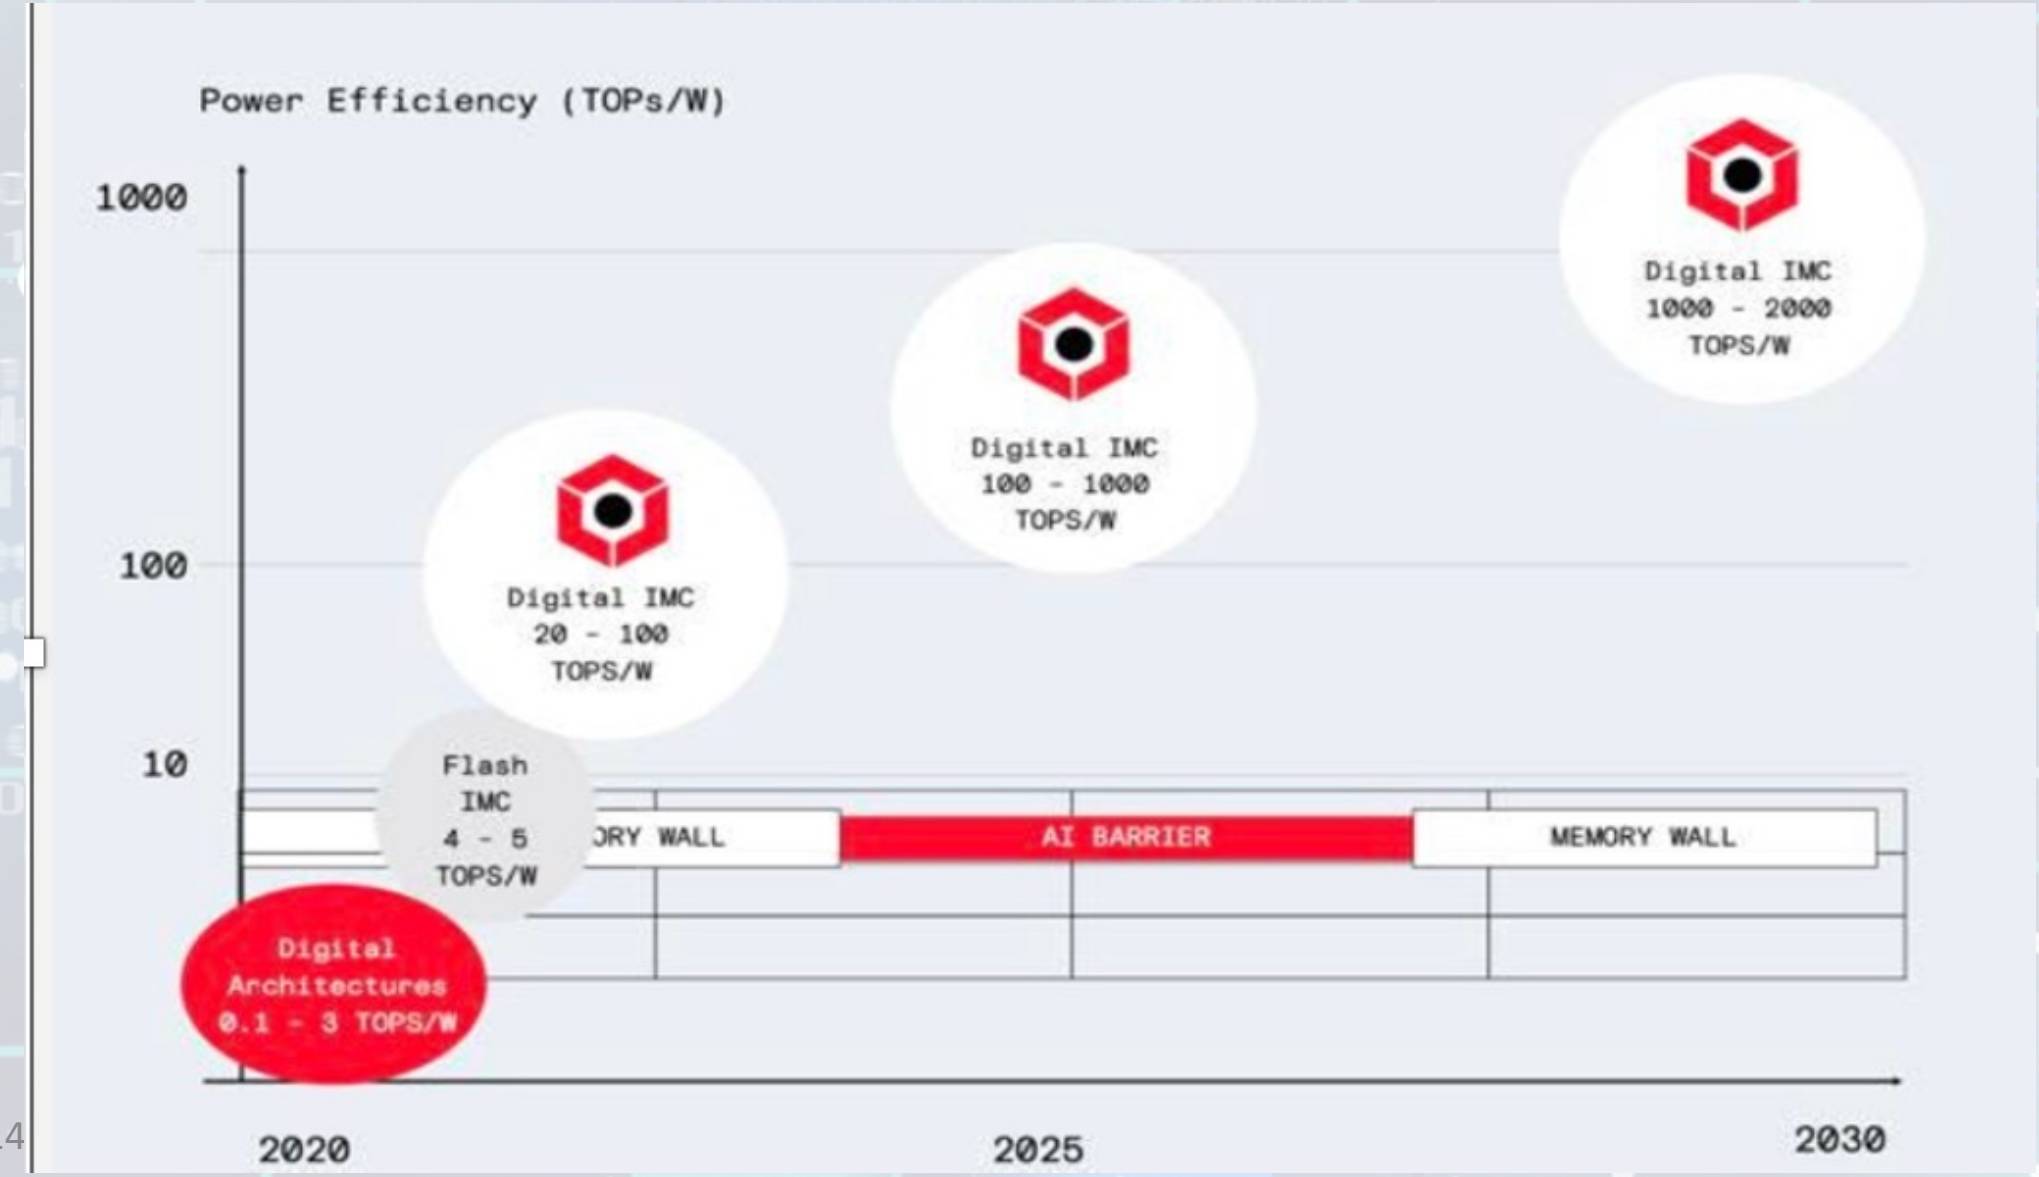

Untether AI: 30 TFLOPS/Watt

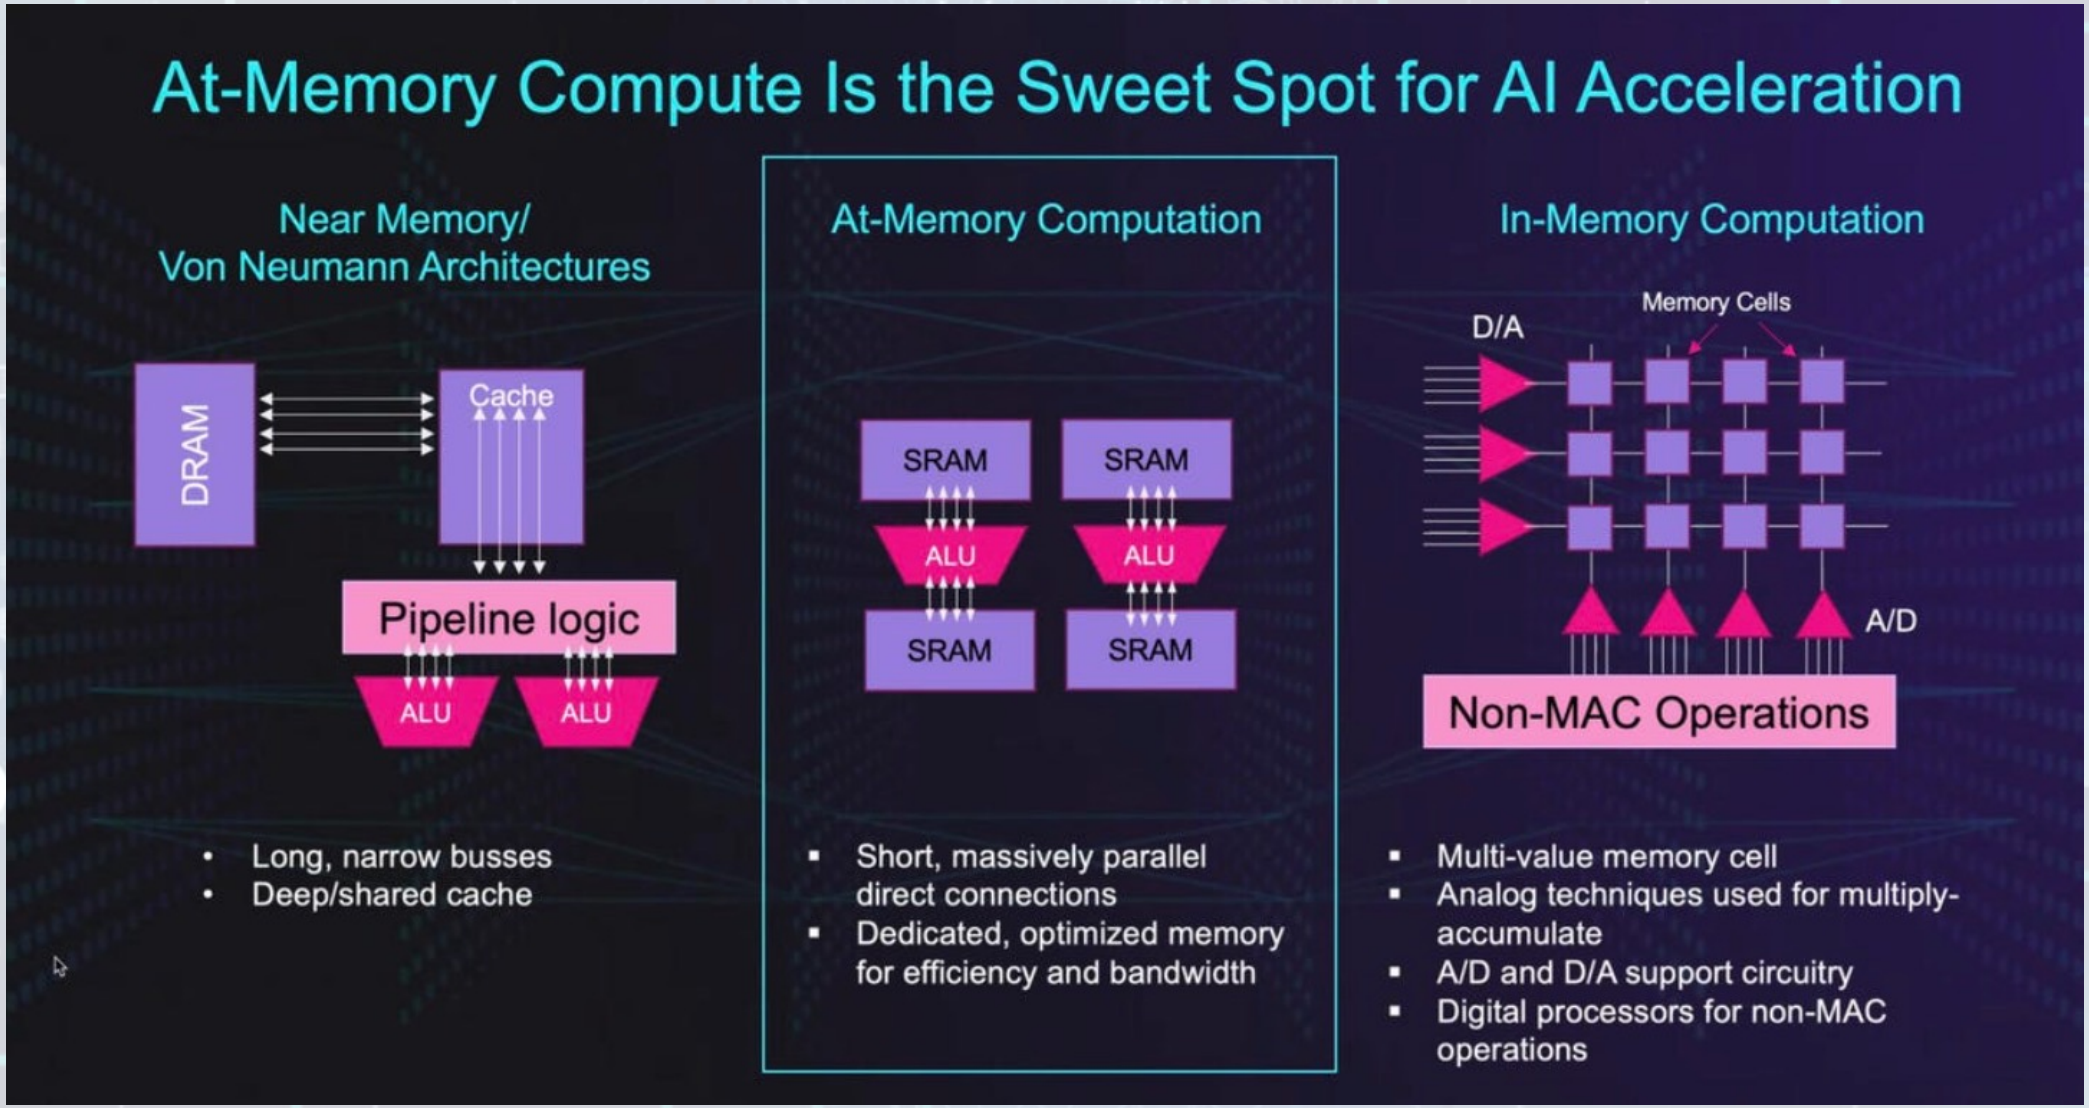

Mythic: Here comes Analog Compute

MYTHIC

Analog - Ideal for Low-Latency Edge AI

- Groundbreaking analog compute-in-memory technology
  - High-density parameter storage with flash
  - Ultra-low latency at very-low power levels
  - Deterministic execution
- Technology and architecture than can easily scale from endpoints to edge-servers

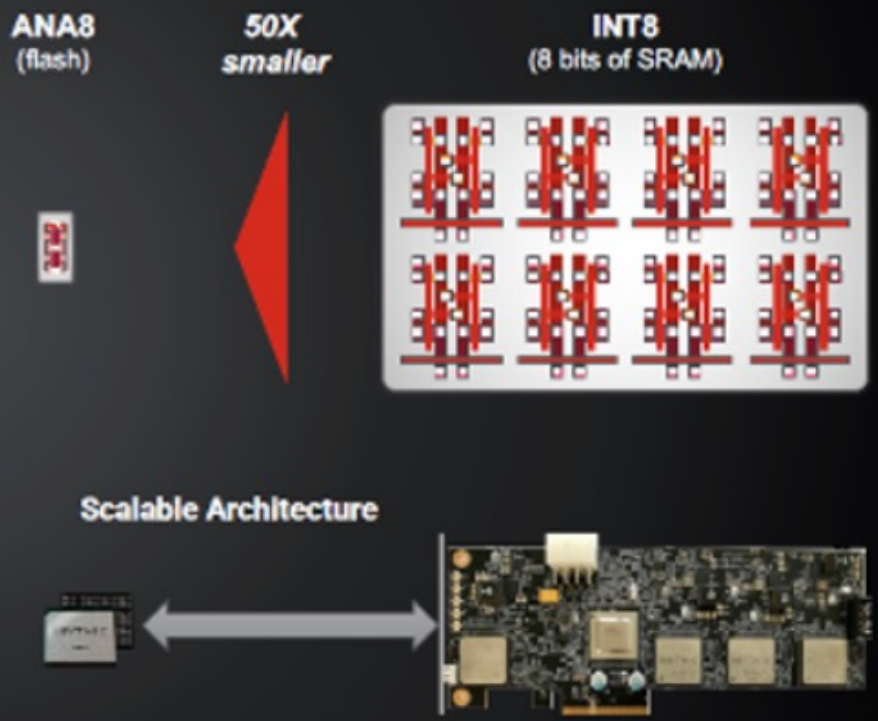

Esperanto: 1000 RISC-V Cores

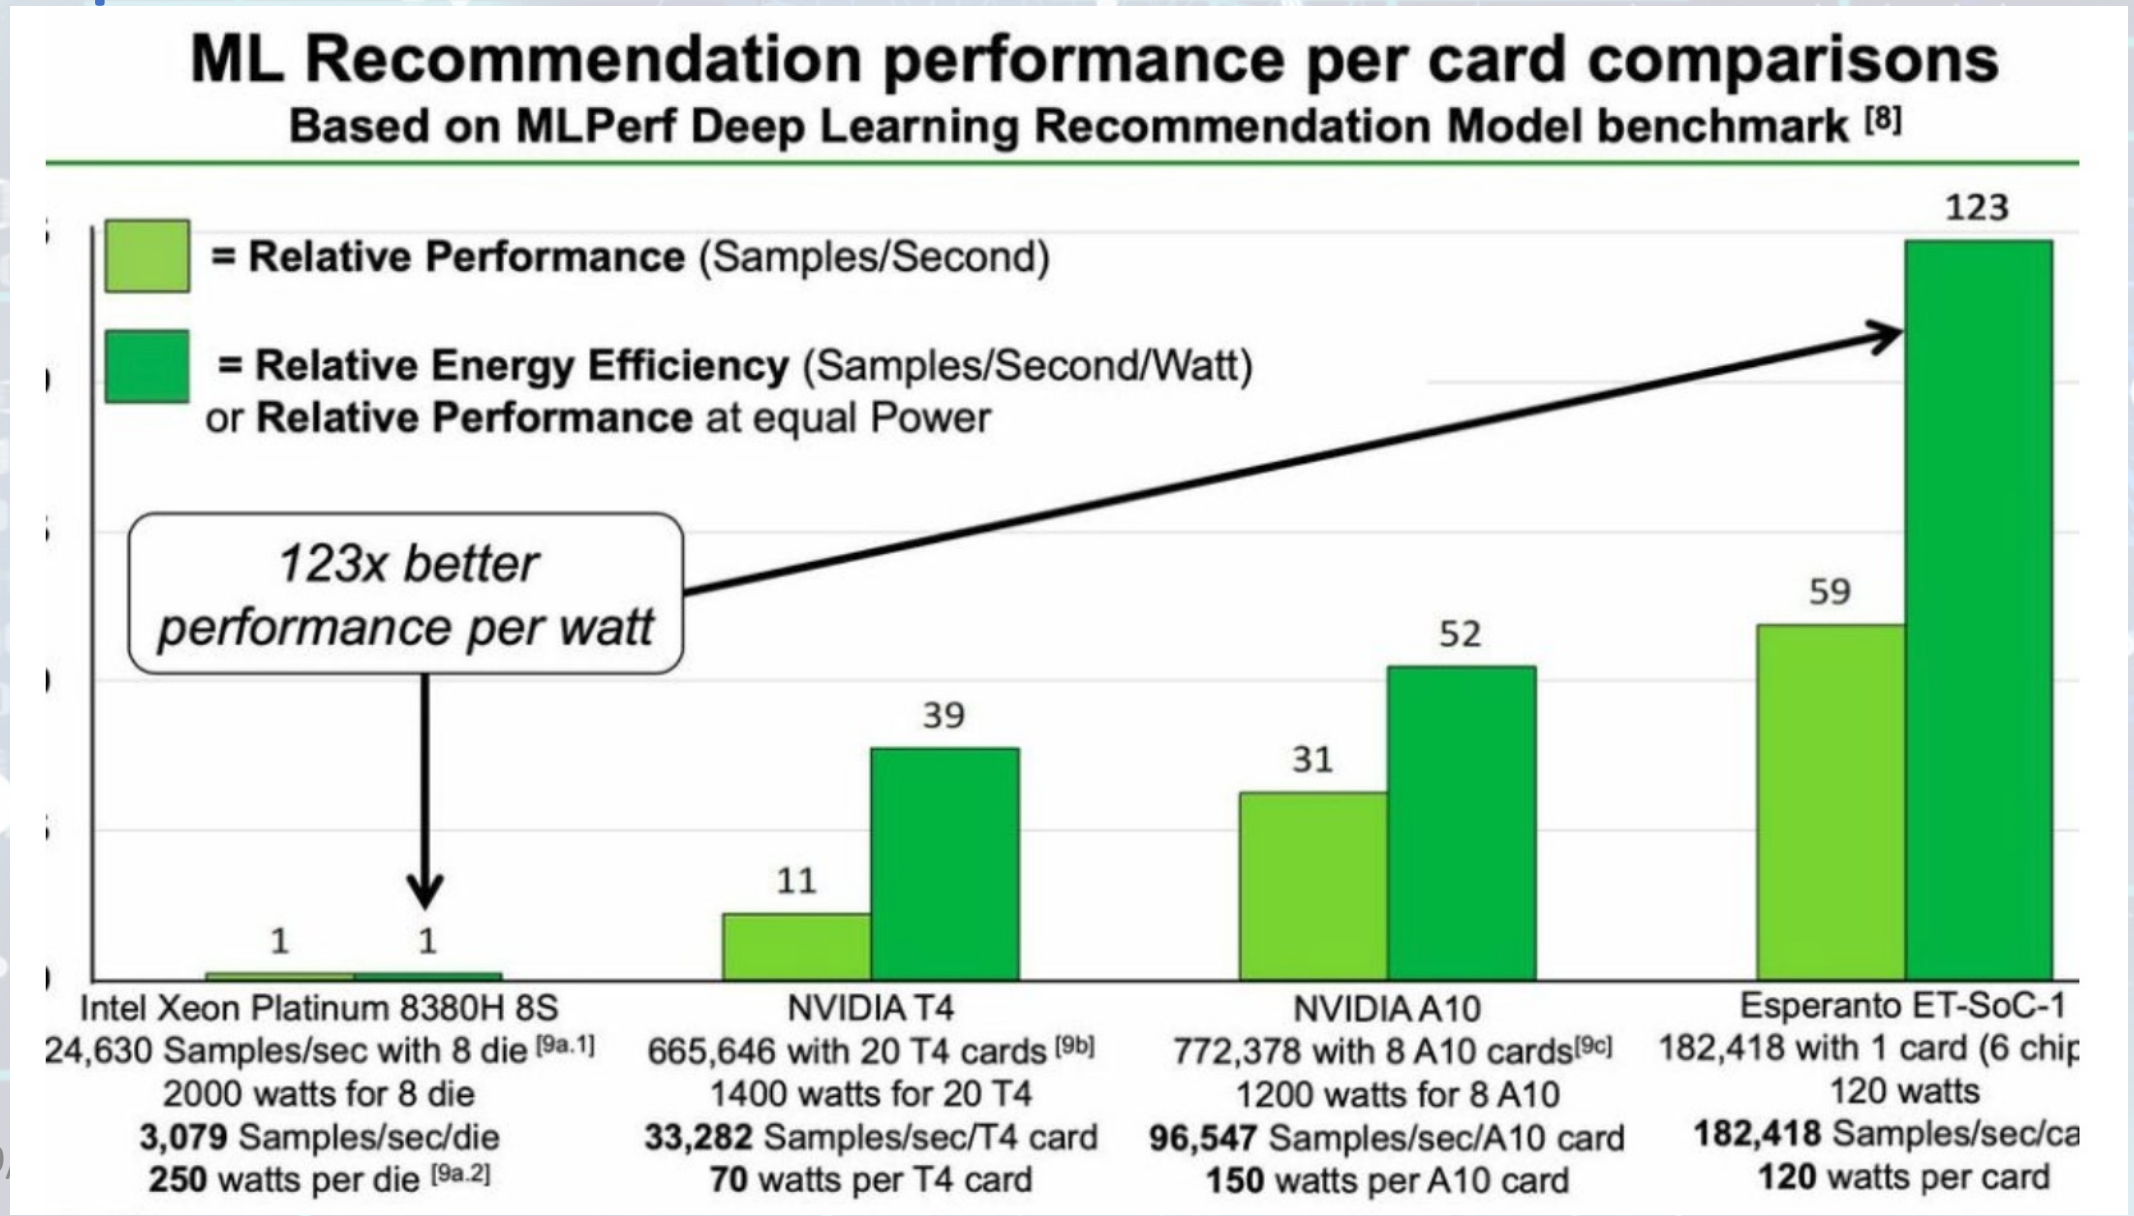

SiMa.ai MLSoC

Machine Learning SoC Device

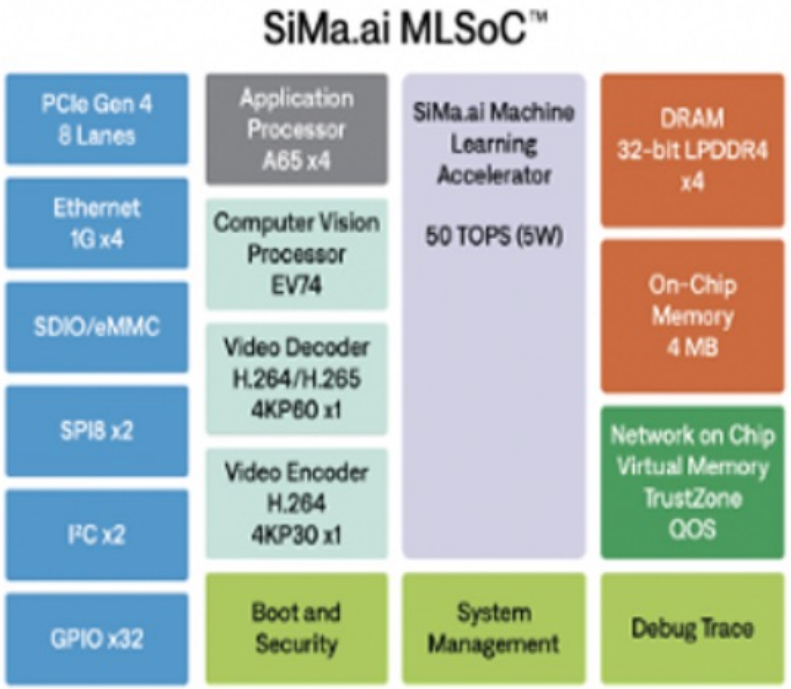

First **software-centric** purpose-built MLSoC platform with **push-button** performance

AMD Instinct MI200

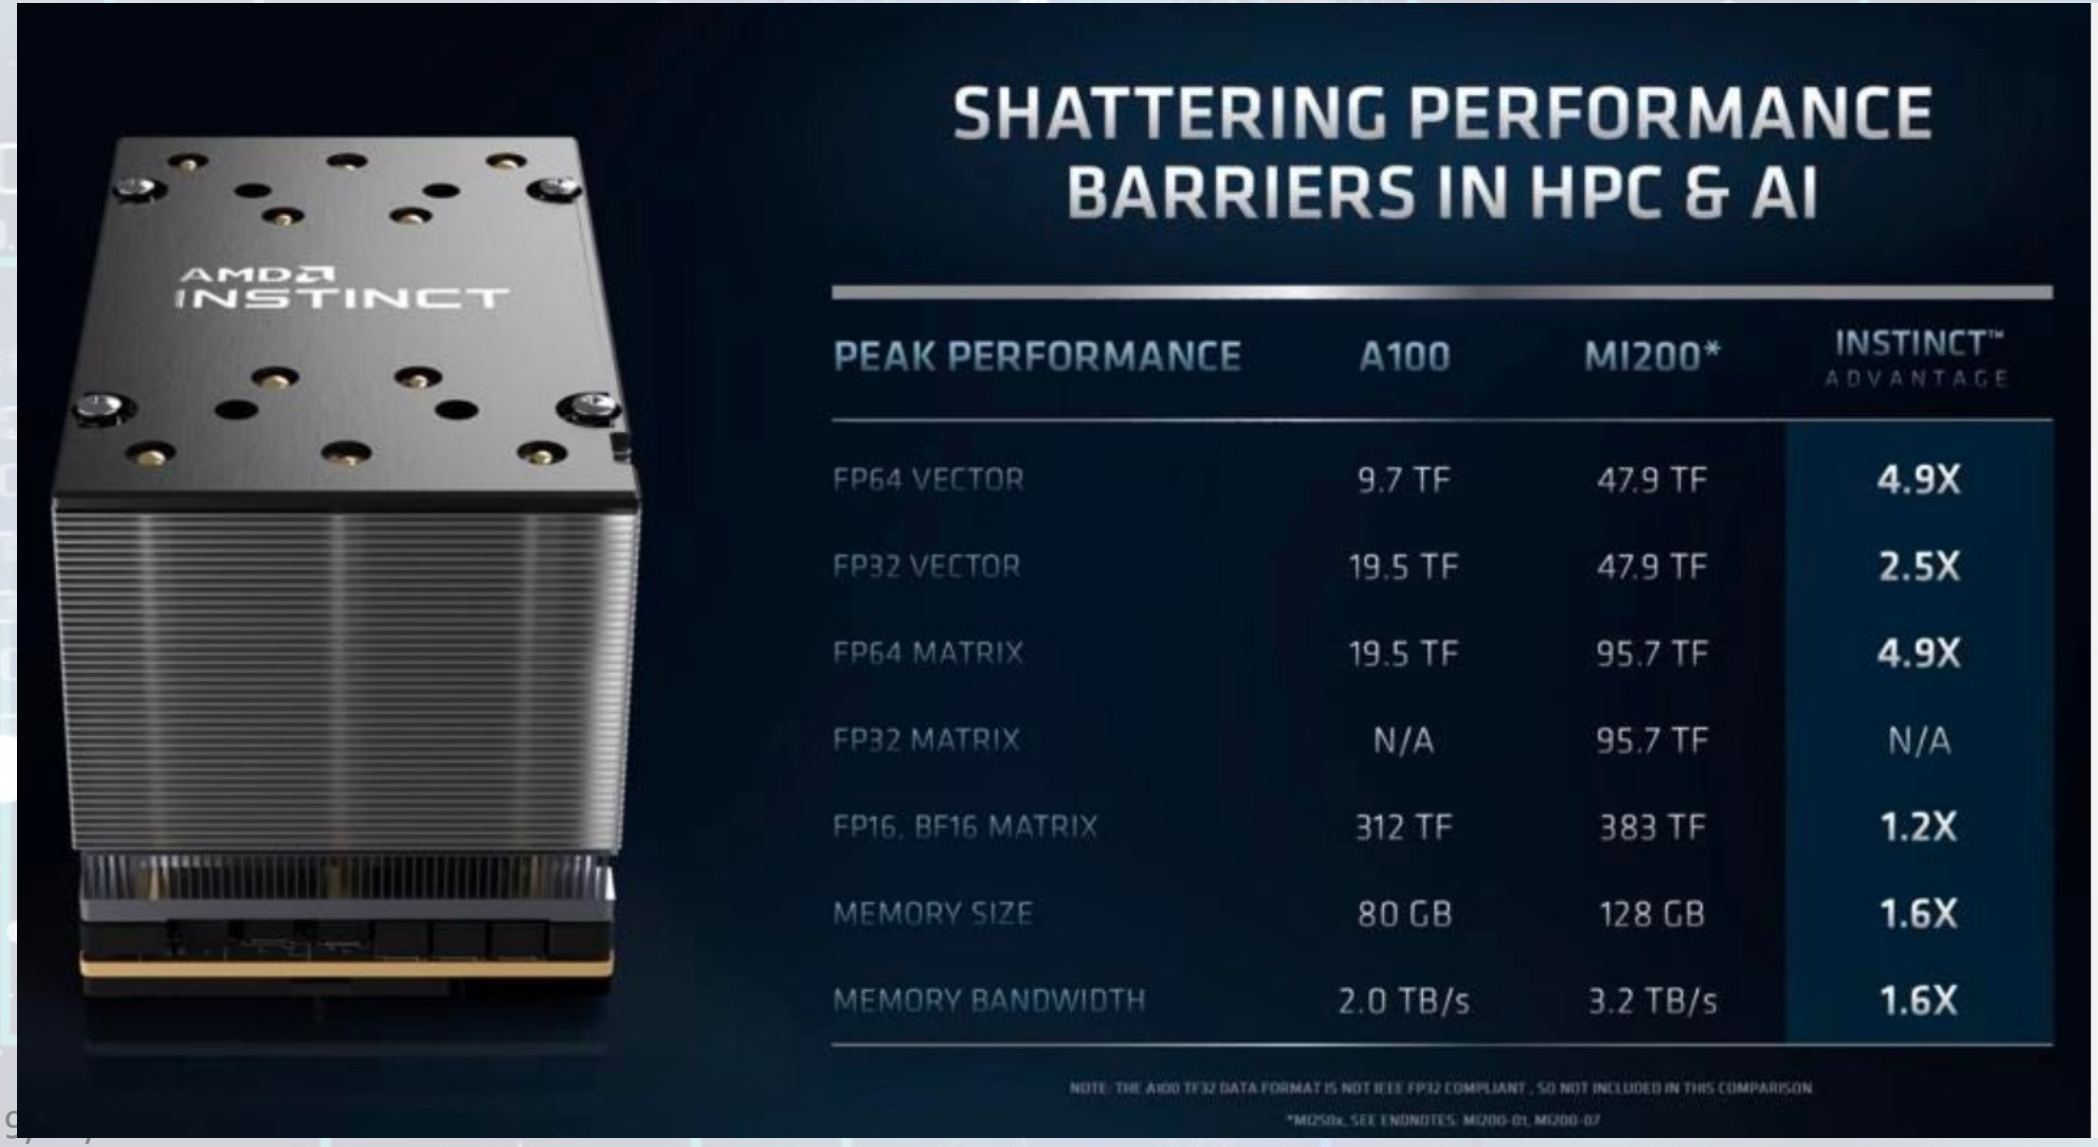

Graphcore: Thinking BIG

Intel Habana Labs Gaudi2: 2X A100

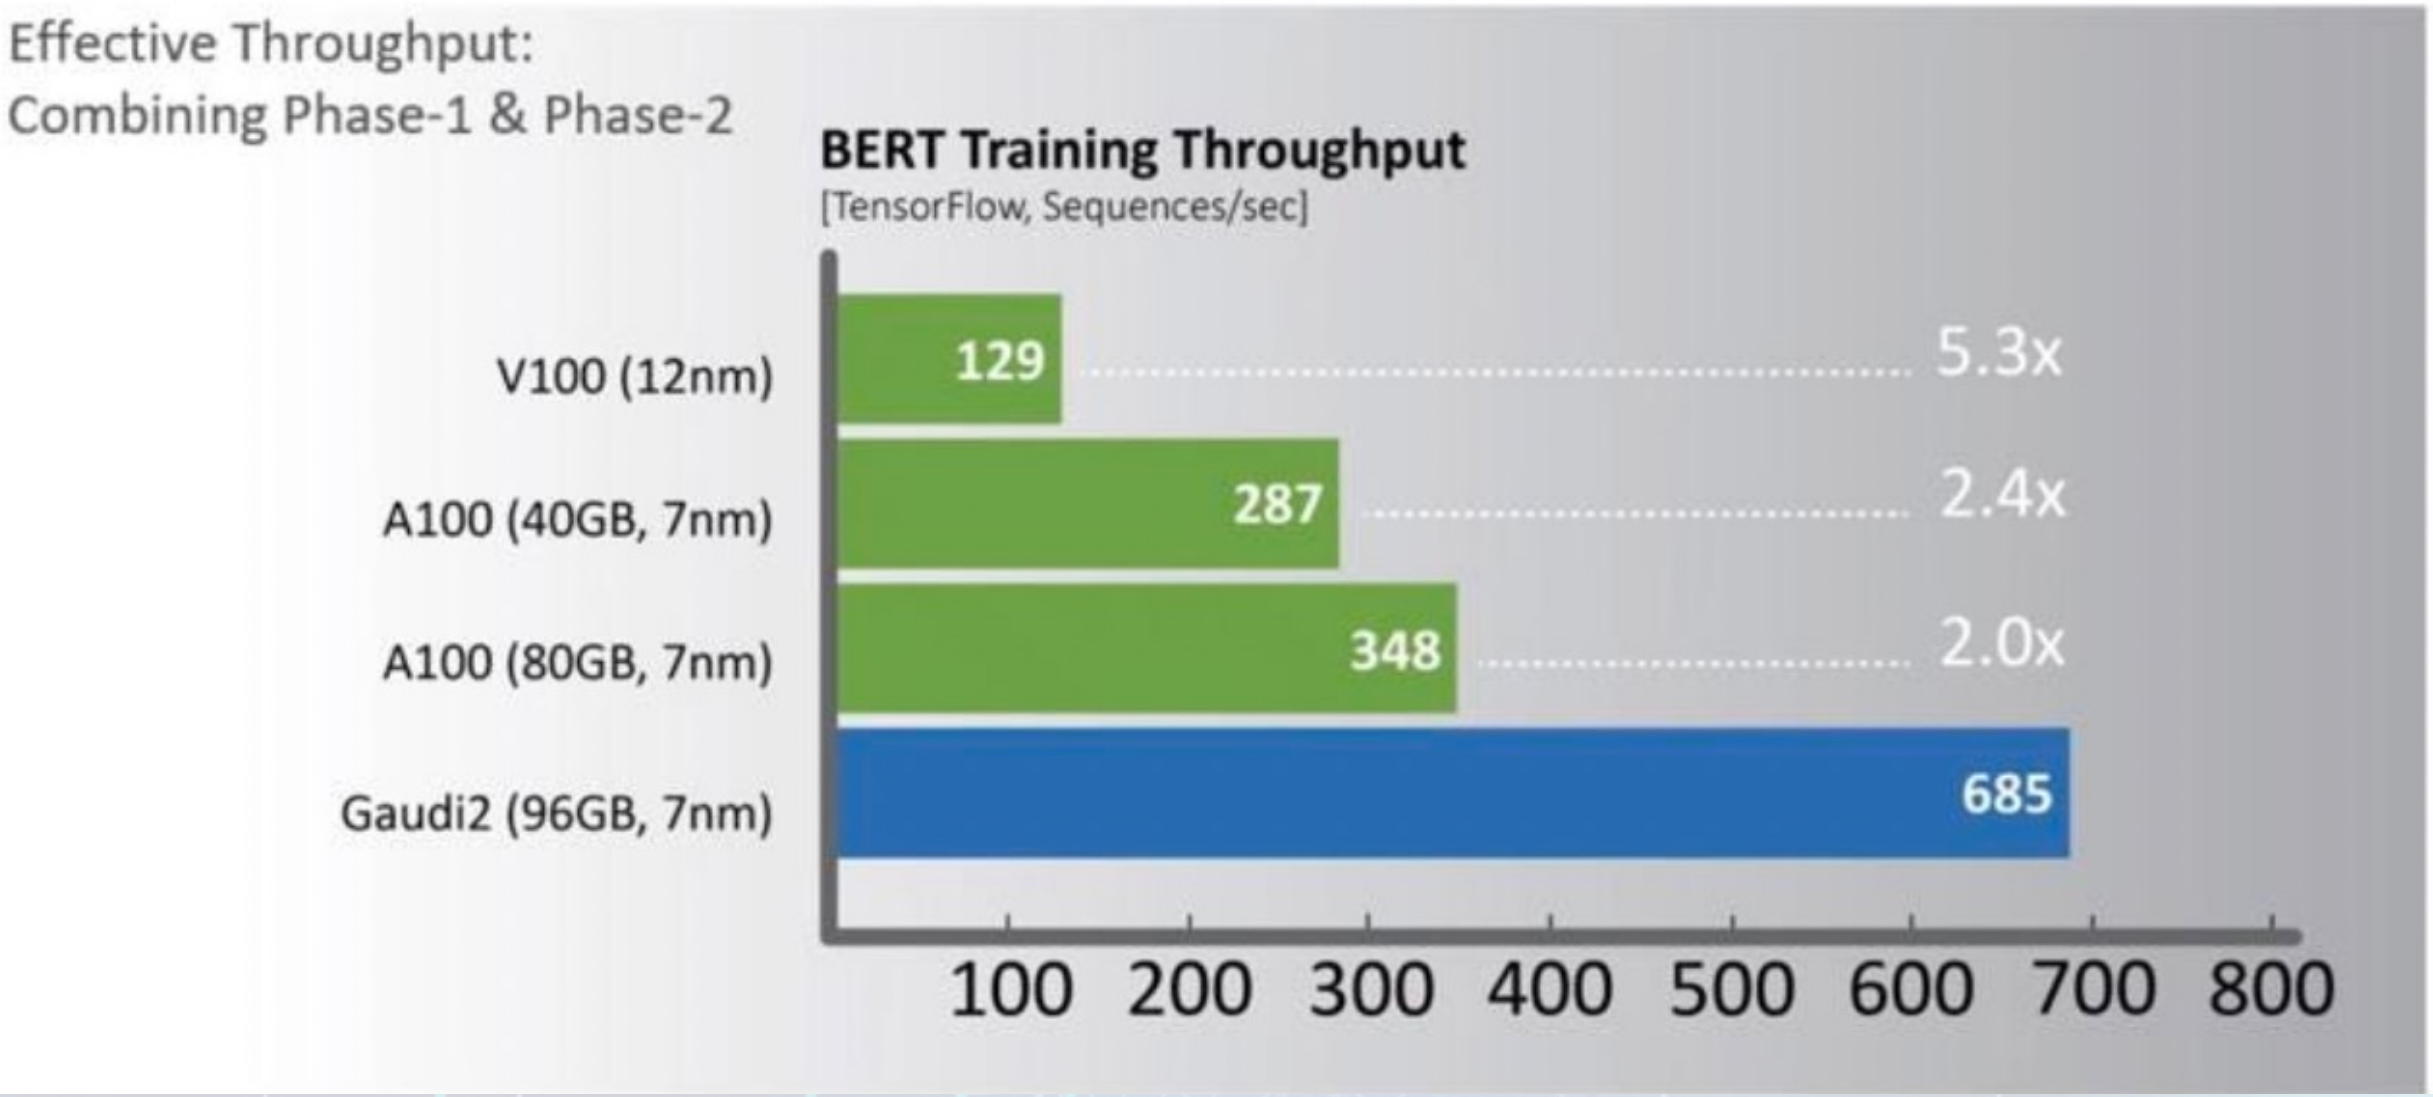

Grace Superchips:
The Future of High Performance

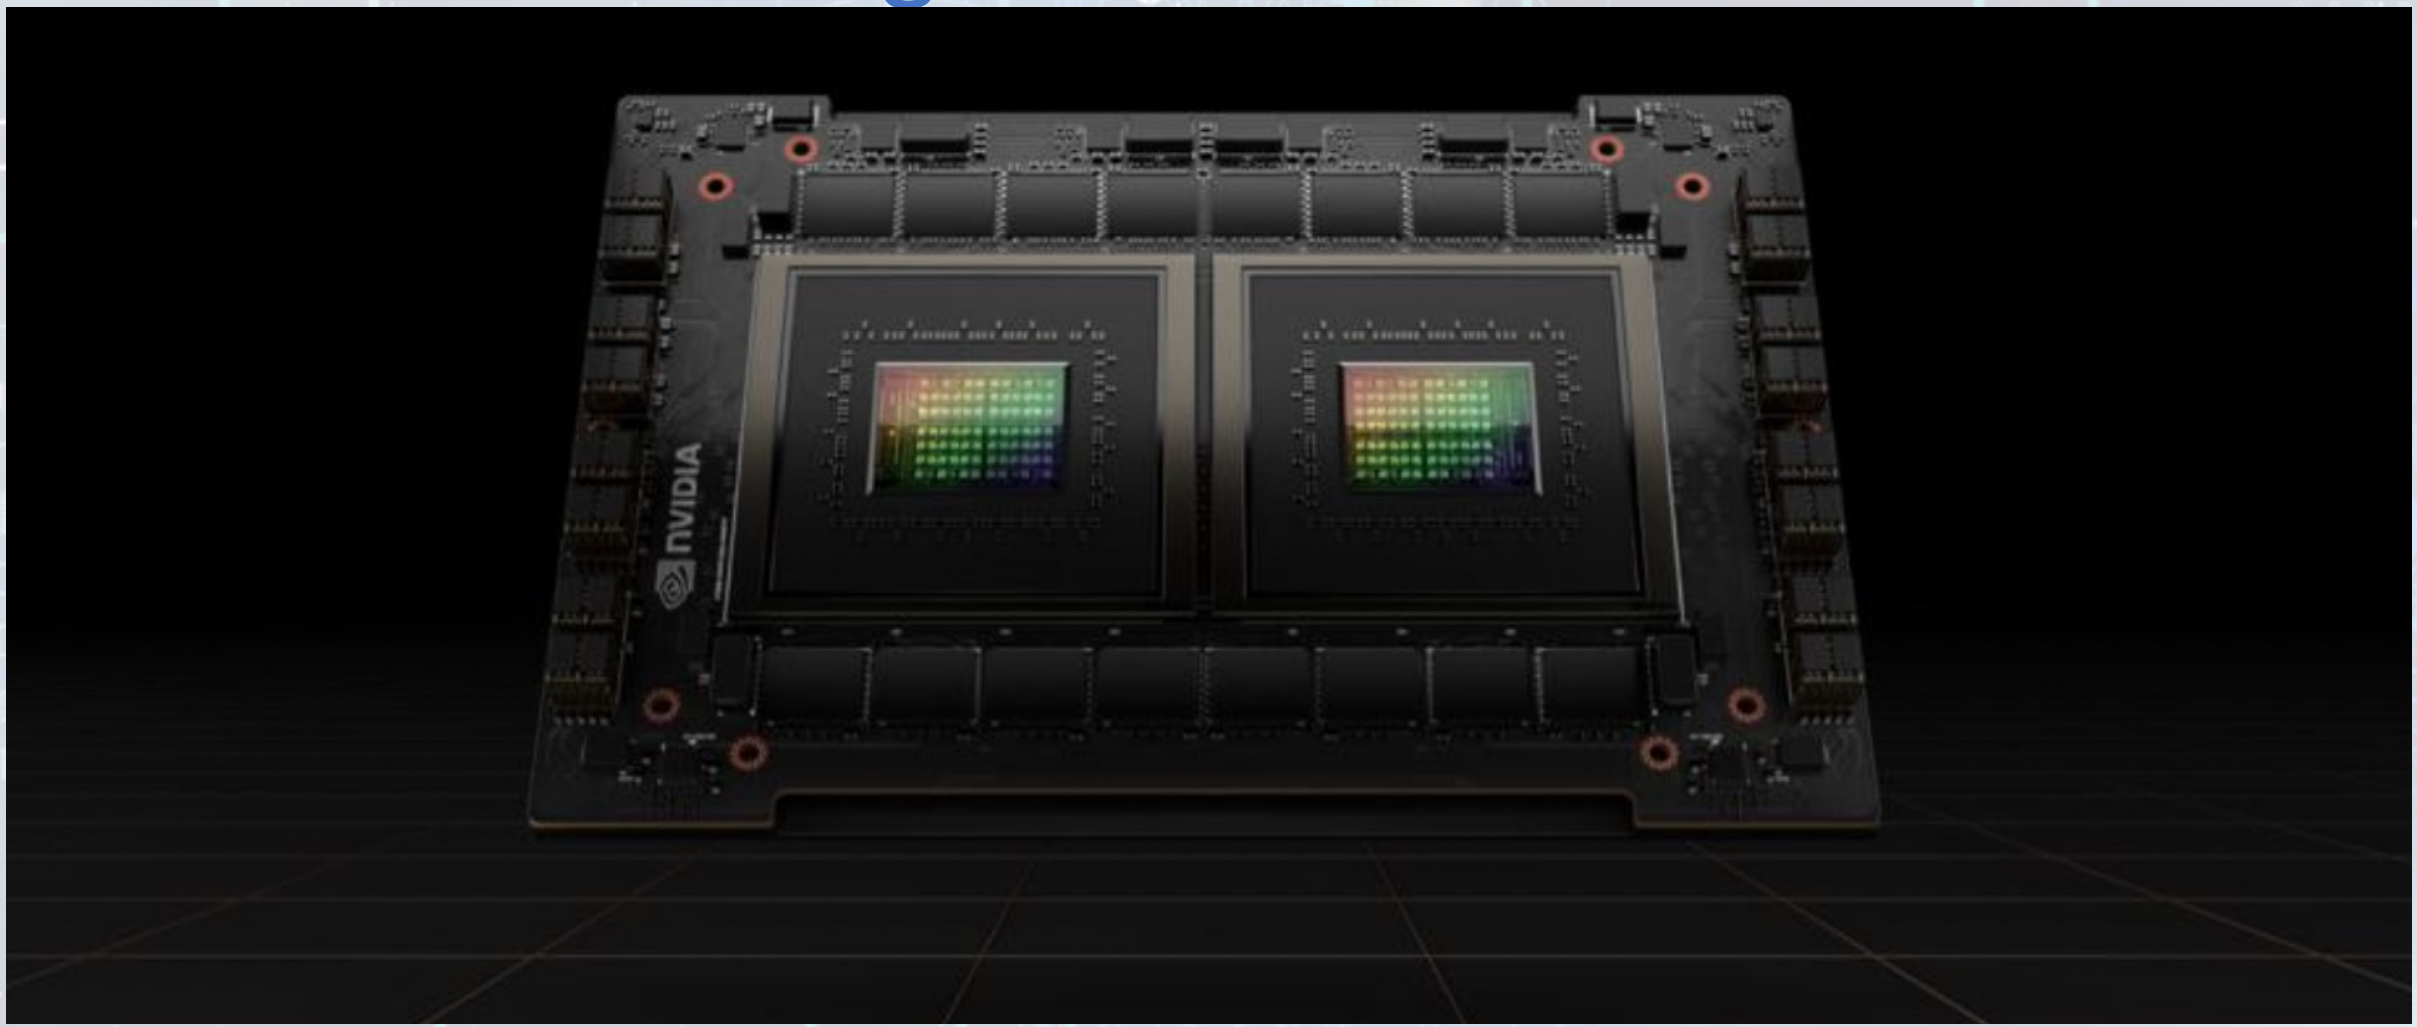

Closing Thoughts

- NVIDIA is practically unassailable in data center training, where Intel, Cerebras, Graphcore, SambaNova, Tenstorrent, and Groq are all attacking.
  - So, if you want to go after this space, you need magic tech and a lot of money
- The Edge market(s) are ripe with opportunities with many niches.
- Software matters even more than you think
  - NVIDIA improved Jetson by **50%** last MLPerf round
- BIG networks (LLMs) will eat the world
  - NVIDIA
  - Graphcore
  - Cerebras

LLMs being used to create images from text

Prompt: "oil on canvas painting + romanticism + landscape + a hay wain pulled by two horses as it crosses a river + a backdrop of mountains, trees, and clouds in the background + simple and idyllic depiction of rural life in England"

Left: The Hay Wain by John Constable (Public Domain). Right: Image by Alberto Romero via Midjourney

THANK YOU!

Have a great conference!

Our clients include:

Come visit our site for news and research reports
http://www.Cambrian-ai.com